Causality and Responsibility for Formal Verification and Beyond

Hana Chockler
Department of Informatics
King's College London
hana.chockler@kcl.ac.uk

The theory of actual causality, defined by Halpern and Pearl, and its quantitative measure – the degree of responsibility – was shown to be extremely useful in various areas of computer science due to a good match between the results it produces and our intuition. In this paper, I describe the applications of causality to formal verification, namely, explanation of counter-examples, refinement of coverage metrics, and symbolic trajectory evaluation. I also briefly discuss recent applications of causality to legal reasoning.

1 Introduction

The definition of causality given by Halpern and Pearl [15], like other definitions of causality in the philosophy literature going back to Hume [19], is based on *counterfactual dependence*. Essentially, event A is a cause of event B if, had A not happened (this is the counterfactual condition, since A did in fact happen) then B would not have happened. Unfortunately, this definition does not capture all the subtleties involved with causality. For example, suppose that Suzy and Billy both pick up rocks and throw them at a bottle (the example is due to Hall [14]). Suzy's rock gets there first, shattering the bottle. Since both throws are perfectly accurate, Billy's would have shattered the bottle had it not been preempted by Suzy's throw. (This story is taken from [14].) Thus, according to the counterfactual condition, Suzy's throw is not a cause for shaterring the bottle. This problem is dealt with in [15] by, roughly speaking, taking A to be a cause of B if B counterfactually depends on A under some contingency. For example, Suzy's throw is a cause of the bottle shattering because the bottle shattering counterfactually depends on Suzy's throw, under the contingency that Billy doesn't throw. It may seem that this solves one problem only to create another. While this allows Suzy's throw to be a cause of the bottle shattering, it also seems to allow Billy's throw to be a cause too, which seems counter-intuitive to most people. As is shown in [17], it is possible to build a more sophisticated model that expresses the subtlety of pre-emption in this case, using auxiliary variables to express the order in which the rocks hit the bottle and preventing Billy's throw from being a cause of the bottle shattering. One moral of this example is that, according to the [17] definitions, whether or not A is a cause of B depends in part on the model used.

Halpern and Pearl's definition of causality, while extending and refining the counter-factual definition, still treats causality as an all-or-nothing concept. That is, A is either a cause of B or it is not. The concept of responsibility, introduced in [5], presents a way to quantify causality and hence the ability to measure the degree of influence (aka "the degree of responsibility") of different causes on the outcome. For example, suppose that Mr. B wins an election against Mr. G. If the vote has been 11–0, it is clear that each of the people who voted for Mr. B is a cause of him winning, but the degree of responsibility of each voter for Mr. B is lower than in a vote 6–5.

Thinking in terms of causality and responsibility was shown to be beneficial for a seemingly unrelated area of research, namely formal verification (model checking) of computerised systems. In *model*

G. Gössler and O. Sokolsky (Eds.): 1st Workshop on Causal Reasoning for Embedded and safety-critical Systems Technologies (CREST'16)
EPTCS 224, 2016, pp. 1–8, doi:10.4204/EPTCS.224.1

© Hana Chockler
This work is licensed under the Creative Commons Attribution License.

checking, we verify the correctness of a finite-state system with respect to a desired behavior by checking whether a labeled state-transition graph that models the system satisfies a specification of this behavior [12]. If the answer to the correctness query is negative, the tool provides a counterexample to the satisfaction of the specification in the system. These counterexamples are used for debugging the system, as they demonstrate an example of an erroneous behaviour [13]. As counterexamples can be very long and complex, there is a need for an additional tool explaining "what went wrong", that is, pinpointing the *causes* of an error on the counterexample. As I describe in more details in the following sections, the causal analysis of counterexamples, described in [2], is an integral part of an industrial hardware verification platform RuleBase of IBM [25].

On the other hand, if the answer is positive, the tools usually perform some type of a sanity check, to verify that the positive result was not caused by an error or underspecification (see [20] for a survey). *Vacuity* check, which is the most common sanity check and is a standard in industrial model checking tools, was first defined by Beer et al. [3] as a situation, where the property passes in a "non-interesting way", that is, a part of a property does not affect the model checking procedure in the system. Beer et al. state that vacuity was a serious problem in verification of hardware designs at IBM: "our experience has shown that typically 20% of specifications pass vacuously during the first formal-verification runs of a new hardware design, and that vacuous passes always point to a real problem in either the design or its specification or environment" [3]. The general vacuity problem was formalised by Kupferman and Vardi, who defined a vacuous pass as a pass, where some subformula of the original property can be replaced by its $\perp$ value without affecting the satisfaction of the property in the system [21], and there is a plethora of subsequent papers and definitions addressing different aspects and nuances of vacuous satisfaction.

Note that vacuity, viewed from the point of view of causal analysis, is *counterfactual causality*. Indeed, a property passes non-vacuously if each its subformula, if replaced by $\perp$, causes falsification of the property in the system. Due to the nature of specifications as more general than implementations, it is quite natural to expect that the specification allows some degree of freedom in how it is going to be satisfied. Consider, for example, a specification $\mathbf{G}(p \vee q)$, meaning that either p or q should hold in each state. If in the system under verification both p and q are **true** in all states, the standard vacuity check will alert the verification engineer to vacuity in p and in q. In contrast, introducing a contingency where p is set to **false** causes the result of model checking to counterfactually depend on q (and similarly for q and p), indicating that both p and q play some role in the satisfaction of the specification in the system.

Coverage check is a concept "borrowed" from testing and simulation-based verification, where various coverage metrics are traditionally used as heuristic measures of exhaustiveness of the verification procedure [29]. In model checking, the suitable coverage concept is *mutation coverage*, introduced by Hoskote et al. [18], and formalized in [8, 7, 6, 9]. In this definition, an element of the system under verification is considered *covered* by the specification if changing (*mutating*) this element falsifies the specification in the system. Note, again, that this definition is, essentially, a counter-factual causality. As a motivating example to the necessity to finer-grained analysis consider the property $\mathbf{F}req$, meaning in every computation there is at least one request. Now consider a system in which requests are sent frequently, resulting in several requests on each of the computational paths. All these requests will be considered not covered by the mutation coverage metric, as changing each one of them separately to **false** does not falsify the property; and yet, each of these requests plays a role in the satisfaction of the specification (or, using the notions from causality, for each request there exists a contingency that creates a counterfactual dependence between this request and the satisfaction of the specification) [11].

While causality allows to extend the concepts of vacuity and coverage to include elements that *affect* the satisfaction of the specification in the system in some way, it is still an all-or-nothing concept. Harnessing the notion of responsibility to measure the influence of different elements on the success or

failure of the model checking process introduces the quantification aspect, providing a finer-grained analysis. While, as I discuss in more detail below, the full-blown responsibility computation is intractable for all but very small systems, introducing a *threshold* on the value of responsibility, in order to detect only the most influential causes, reduces the complexity and makes the computation manageable for real systems [11].

The quantification provided by the notion of responsibility and the distinction between influential and non-influential causes have been applied to the symbolic trajectory evaluation, where ordering the causes by their degree of responsibility was demonstrated to be a good heuristic for instantiating a minimal number of variables that is sufficient to determine the output value of the circuit [4].

In the next sections I provide a brief overview of the relevant concepts and describe the applications of causality to formal verification in more detail.

2 Causality and Responsibility – Definitions

In this section, I briefly review the details of Halpern and Pearl's definition (HP) of causal models and causality [15] and the definitions of responsibility and blame [5] in causal models.

2.1 Causal models

A *signature* is a tuple $\mathbf{S} = \langle \mathbf{U}, \mathbf{V}, \mathbf{R} \rangle$, where $\mathbf{U}$ is a finite set of *exogenous* variables, $\mathbf{V}$ is a finite set of *endogenous* variables, and $\mathbf{R}$ associates with every variable $Y \in \mathbf{U} \cup \mathbf{V}$ a finite nonempty set $\mathbf{R}(Y)$ of possible values for Y. Intuitively, the exogenous variables are ones whose values are determined by factors outside the model, while the endogenous variables are ones whose values are ultimately determined by the exogenous variables. A (recursive) *causal model* over signature $\mathbf{S}$ is a tuple $M = \langle \mathbf{S}, \mathbf{F} \rangle$, where $\mathbf{F}$ associates with every endogenous variable $X \in \mathbf{V}$ a function F_X such that $F_X : (\times_{U \in \mathbf{U}} \mathbf{R}(U) \times (\times_{Y \in \mathbf{V} \setminus \{X\}} \mathbf{R}(Y))) \rightarrow \mathbf{R}(X)$, and functions have no circular dependency. That is, F_X describes how the value of the endogenous variable X is determined by the values of all other variables in $\mathbf{U} \cup \mathbf{V}$. If all variables have only two values, we say that M is a *binary causal model*.

We can describe (some salient features of) a causal model M using a *causal network*, which is a graph with nodes corresponding to the variables in M and edges corresponding to the dependence between variables. We focus our attention on *recursive models* – those whose associated causal network is a directed acyclic graph.

A causal network is a graph with nodes corresponding to the random variables in $\mathbf{V}$ and an edge from a node labeled X to one labeled Y if F_Y depends on the value of X. Intuitively, variables can have a causal effect only on their descendants in the causal network; if Y is not a descendant of X, then a change in the value of X has no affect on the value of Y. A setting $\vec{u}$ for the variables in $\mathbf{U}$ is called a *context*. It should be clear that if M is a recursive causal model, then there is always a unique solution to the equations in M, given a context.

Given a causal model $M = (\mathbf{S}, \mathbf{F})$, a (possibly empty) vector $\vec{X}$ of variables in $\mathbf{V}$, and a vector $\vec{x}$ of values for the variables in $\vec{X}$, a new causal model, denoted $M_{\vec{X} \leftarrow \vec{x}}$, is defined as identical to M, except that the equation for the variables $\vec{X}$ in $\mathbf{F}$ is replaced by $\vec{X} = \vec{x}$. Intuitively, this is the causal model that results when the variables in $\vec{X}$ are set to $\vec{x}$ by some external action that affects only the variables in $\vec{X}$ (and overrides the effects of the causal equations).

A causal formula φ – a Boolean combination of events capturing the value of variables in the model – is **true** or **false** in a causal model, given a *context*. We write $(M, \vec{u}) \models \varphi$ if φ is **true** in the causal model

M, given the context $\vec{u}$. $(M,\vec{u}) \models [\vec{Y} \leftarrow \vec{y}](X = x)$ if the variable X has value x in the unique solution to the equations in $M_{\vec{Y}\leftarrow\vec{y}}$ in context $\vec{u}$. We now review the HP definition of causality..

Definition 2.1 (Cause [15]) $\vec{X} = \vec{x}$ *is a* cause *of* φ *in* $(M,\vec{u})$ *if the following three conditions hold:*

AC1. $(M,\vec{u}) \models (\vec{X} = \vec{x}) \wedge \varphi$.

AC2. *There exist a partition* $(\vec{Z},\vec{W})$ *of* $\mathbf{V}$ *with* $\vec{X} \subseteq \vec{Z}$ *and some setting* $(\vec{x}',\vec{w})$ *of the variables in* $(\vec{X},\vec{W})$ *such that if* $(M,\vec{u}) \models Z = z^*$ *for* $Z \in \vec{Z}$*, then*

(a) $(M,\vec{u}) \models [\vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w}]\neg\varphi$.

(b) $(M,\vec{u}) \models [\vec{X} \leftarrow \vec{x}, \vec{W}' \leftarrow \vec{w}, \vec{Z}' \leftarrow \vec{z}^*]\varphi$ *for all subsets* $\vec{Z}'$ *of* $\vec{Z} \setminus \vec{X}$ *and all subsets* $\vec{W}'$ *of* $\vec{W}$*. The tuple* $(\vec{W},\vec{w},\vec{x}')$ *is said to be a* witness *to the fact that* $\vec{X} = \vec{x}$ *is a cause of* φ.

AC3. $(\vec{X} = \vec{x})$ *is minimal; no subset of* $\vec{X}$ *satisfies AC2.*

Essentially, Definition 2.1 extends the counterfactual definition of causality by considering contingencies $\vec{W}$ – changes in the current context that by themselves do not change the value of φ, but create a counterfactual dependence between the value of a $\vec{X}$ and the value of φ. The variables in $\vec{Z}$ should be thought of as describing the "active causal process" from X to φ, and are needed in order to make sure that, while introducing contingencies, we preserve the causal process from $\vec{X}$ to φ. The minimality requirement AC3 is needed in order to avoid adding irrelevant variables to $\vec{X}$.

We note that Halpern recently updated the definition of causality, changing the concept to focus on the variables that are *frozen* in their original values, rather than considering contingencies [16]. Since the existing work on the applications of causality to formal verification uses the previous definition, we continue using it in this paper.

2.2 Responsibility and Blame

Causality is a "0–1" concept; $\vec{X} = \vec{x}$ is either a cause of φ or it is not. Now consider two voting scenarios: in the first, Mr. G beats Mr. B by a vote of 11–0. In the second, Mr. G beats Mr. B by a vote of 6–5. According to the HP definition, all the people who voted for Mr. G are causes of him winning. While this does not seem so unreasonable, it does not capture the intuition that each voter for Mr. G is more critical to the victory in the case of the 6–5 vote than in the case of the 11–0 vote. The notion of *degree of responsibility*, introduced by Chockler and Halpern [5], extends the notion of causality to capture the differences in the degree of criticality of causes. In the case of the 6–5 vote, no changes have to be made to make each voter for Mr. G critical for Mr. G's victory; if he had not voted for Mr. G, Mr. G would not have won. Thus, each voter has degree of responsibility 1 (i.e., $k = 0$). On the other hand, in the case of the 11–0 vote, for a particular voter to be critical, five other voters have to switch their votes; thus, $k = 5$, and each voter's degree of responsibility is $1/6$. This notion of degree of responsibility has been shown to capture (at a qualitative level) the way people allocate responsibility [22].

Definition 2.2 (Degree of Responsibility [5]) *The* degree of responsibility of $\vec{X} = \vec{x}$ for φ in $(M,\vec{u})$*, denoted* $\mathrm{dr}((M,\vec{u}),(\vec{X} \leftarrow \vec{x}),\varphi)$*, is* 0 *if* $\vec{X} = \vec{x}$ *is not a cause of* φ *in* $(M,\vec{u})$*; it is* $1/(k+1)$ *if* $\vec{X} = \vec{x}$ *is a cause of* φ *in* $(M,\vec{u})$ *according to Definition 2.1 with* $\vec{W}$ *of size k being a smallest witness to the fact that* $\vec{X} = \vec{x}$ *is a cause of* φ.

When determining responsibility, it is assumed that everything relevant about the facts of the world and how the world works (which we characterize in terms of *structural equations*) is known. But this misses out on important component of determining what Chockler and Halpern call *blame*: the epistemic state. Formally, the degree of blame, introduced by Chockler and Halpern is the expected degree of

responsibility [5]. This is perhaps best understood by considering a firing squad with ten excellent marksmen (the example is by Tim Williamson). Only one of them has live bullets in his rifle; the rest have blanks. The marksmen do not know which of them has the live bullets. The marksmen shoot at the prisoner and he dies. The only marksman that is the cause of the prisoner's death is the one with the live bullets. That marksman has degree of responsibility 1 for the death; all the rest have degree of responsibility 0. However, each of the marksmen has degree of blame $1/10$.

An agent's uncertainty is modelled by a pair $(K, \Pr)$, where K is a set of pairs of the form $(M, \vec{u})$, where M is a causal model and $\vec{u}$ is a context, and Pr is a probability distribution over K. Note that probability is used here in a rather non-traditional sense, to capture the epistemic state of an agent, rather than an actual probability over values of variables.

Definition 2.3 (Blame [5]) *The* degree of blame of setting $\vec{X}$ to $\vec{x}$ for φ relative to epistemic state $(K, \Pr)$, *denoted* $\mathrm{db}(K, \Pr, \vec{X} \leftarrow \vec{x}, \varphi)$, *is defined as an* expected value *of the degree of responsibility over the probability space* $(K, \Pr)$.

3 Coverage in the framework of causality

The following definition of coverage is based on the study of *mutant coverage* in simulation-based verification [23, 24, 1], and is the one that is adopted in all (or almost all) papers on coverage metrics in formal verification today (see, for example, [18, 8, 7, 6, 9]). For a Kripke structure K, an atomic proposition q, and a state w, we denote by $\tilde{K}_{w,q}$ the Kripke structure obtained from K by flipping the value of q in w.

Definition 3.1 (Coverage) *Consider a Kripke structure K, a specification φ that is satisfied in K, and an atomic proposition $q \in AP$. A state w of K is* q-covered by φ *if $\tilde{K}_{w,q}$ does not satisfy φ.*

It is easy to see that coverage corresponds to the simple counterfactual-dependence approach to causality. Indeed, a state w of K is *q-covered by φ* if φ holds in K and if q had other value in w, then φ would not have been true in K. The following example illustrates the notion of coverage and shows that the counter-factual approach to coverage misses some important insights in how the system satisfies the specification. Let K be a Kripke structure presented with one path, where one request is followed by three grants in subsequent states, and let $\varphi = \mathbf{G}(req \rightarrow \mathbf{F}grant)$ (every request is eventually granted). It is easy to see that K satisfies φ, but that none of the states are covered with respect to *grant*, as flipping the value of *grant* in one of them does not falsify φ in K. On the other hand, representing the model checking procedure as a causal model with a context corresponding to the actual values of atomic propositions in states (see [11] for a formal description of this representation), demonstrates that for each state there exists a contingency where the result counterfactually depends on the value of *grant* in this state; the contingency is removing grants in the two other states. Hence, while none of the states is covered with respect to *grant*, they are all causes of φ in K with the responsibility $1/3$.

In the example above, and typically in the applications of causality to formal verification, there are no structural equations over the internal endogenous variables. However, if we want to express the temporal characteristics of our model – for example, to say that the first *grant* is important, whereas subsequent ones are not – the way to do so is by introducing internal auxiliary variables, expressing the order between the grants in the system.

4 Explanation of Counterexamples Using Causality

Explanation of counterexamples addresses a basic aspect of understanding a counterexample: the task of finding the failure in the trace itself. To motivate this approach, consider a verification engineer, who

is formally verifying a hardware design written by a logic designer. The verification engineer writes a specification – a temporal logic formula – and runs a model checker, in order to check the formula on the design. If the formula fails on the design-under-test (DUT), a counterexample trace is produced and displayed in a trace viewer. The verification engineer does not attempt to debug the DUT implementation (since that is the responsibility of the the logic designer who wrote it). Her goal is to look for some basic information about the manner in which the formula fails on the specific trace. If the formula is a complex combination of several conditions, she needs to know which of these conditions has failed. These basic questions are prerequisites to deeper investigations of the failure. Ben-David et al. present a method and a tool for explaining the trace, without involving the model from which it was extracted [2]. This gives the approach the advantage of being light-weight, as its running time depends only on the size of the specification and the counterexample, which is much smaller than the size of the system. An additional advantage of the tool is that it is independent on the verification procedure, and can be added as an external layer to any tool, or even applied as an explanation of simulation traces. The main idea of the algorithm is to represent the trace and the property that fails on this trace as a causal model and context (see Section 3 on the description of the transformation). Then, the values of signals in specific cycles are viewed as variables, and the set of causes for failure is marked as red dots on the trace that is shown to the user graphically, in form of a timing diagram. The tool is a part of the IBM RuleBase verification platform [25] and is used extensively by the users. While the trace is small compared to the system, the complexity of the exact algorithm for computing causality leads to time-consuming computations; in order to keep the interactive nature of the tool, the algorithm operates in one pass on the trace and computes an *approximate* set of causes, which coincides with the actual set of causes on all but contrived examples. The reader is referred to [2] for more details of the implementation.

5 Responsibility in Symbolic Trajectory Evaluation

Symbolic Trajectory Evaluation (STE) [27] is a powerful model checking technique for hardware verification, which combines symbolic simulation with 3-valued abstraction. Consider a circuit M, described as a Directed Acyclic Graph of nodes that represent gates and latches. For such a circuit, an STE assertion is of the form $A \rightarrow C$, where the *Antecedent* A imposes constraints over nodes of M at different times, and the *Consequent* C imposes requirements on M's nodes at different times. The antecedent may introduce symbolic Boolean variables on some of the nodes. The nodes that are not restricted by A are initialized by STE to the value X (”unknown”), thus obtaining an *abstraction* of the checked model.

STE is successfully used in the hardware industry for verifying very large models with wide data paths [28, 26, 30]. The common method for performing STE is by representing the values of each node in the circuit by *Binary Decision Diagrams (BDDs)*. To avoid the potential state explosion resulting from instantiating all unconstrained nodes, typically the circuit is refined manually in iterations, until the value of the circuit is determined.

To avoid the need for manual refinement (which requires a close familiarity with the structure of the circuit), [4] suggest to compute an approximation of the *degree of responsibility* of each node in the value of the output circuit. Then, the instantiation can proceed in the order of decreasing degree of responsibility. The idea behind this algorithm is that the nodes with the highest degree of responsibility are more likely to influence the value of the circuit, and hence we will avoid instantiating too many nodes. The algorithm was implemented in Intel STE framework for hardware verification and demonstrated better results than manual refinement [4].

6 ...and Beyond

In formal verification, there is no natural application for the notion of blame (Def. 2.3), since the model and the property are assumed to be known. On the other hand, in legal applications it is quite natural to talk about an epistemic state of an agent, representing what the agent knew or should have known. As [15] points out, the legal system does not agree with the structural definitions of causality, responsibility, and blame. However, it is still possible to apply our methodology of representing a problem using causal models in order to improve our understanding and analysis of particular situations. In [10], we make the case of using the framework of actual causality in order to guide legal inquiry. In fact, the concepts of responsibility and blame fit the procedure of legal inquiry very well, since we can capture both the limited knowledge of the participants in the case, and the unknown factors in the case. We use the case of baby P. – a baby that died from continuous neglect and abuse, which was completely missed by the social services and the doctor who examined him – to demonstrate how we can capture the known and unknown factors in the case and attempt to quantify the blame of different parties involved in the case.

References

[1] P. Ammann & P.E. Black (2001): *A specification-based coverage metric to evaluate test sets*. *International Journal of Quality, Reliability and Safety Engineering* 8(4), pp. 275–300, doi:10.1142/S0218539301000530.

[2] I. Beer, S. Ben-David, H. Chockler, A. Orni & R. J. Trefler (2012): *Explaining counterexamples using causality*. *Formal Methods in System Design* 40(1), pp. 20–40. doi:10.1007/s10703-011-0132-2.

[3] I. Beer, S. Ben-David, C. Eisner & Y. Rodeh (2001): *Efficient detection of vacuity in ACTL formulas*. *Formal Methods in System Design* 18(2), pp. 141–162, doi:10.1023/A:1008779610539.

[4] H. Chockler, O. Grumberg & A. Yadgar (2008): *Efficient automatic STE refinement using responsibility*. In: *Proc. 14th Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pp. 233–248. doi:10.1007/978-3-540-78800-3_17.

[5] H. Chockler & J.Y. Halpern (2004): *Responsibility and blame: a structural-model approach*. *Journal of Artificial Intelligence Research (JAIR)* 22, pp. 93–115.

[6] H. Chockler & O. Kupferman (2002): *Coverage of Implementations by Simulating Specifications*. In R.A. Baeza-Yates, U. Montanari & N. Santoro, editors: *Proceedings of 2nd IFIP International Conference on Theoretical Computer Science*, *IFIP Conference Proceedings* 223, Kluwer Academic Publishers, Montreal, Canada, pp. 409–421.

[7] H. Chockler, O. Kupferman, R.P. Kurshan & M.Y. Vardi (2001): *A Practical Approach to Coverage in Model Checking*. In: *Computer Aided Verification, Proc. 13th International Conference*, *Lecture Notes in Computer Science* 2102, Springer-Verlag, pp. 66–78, doi:10.1007/3-540-44585-4_7.

[8] H. Chockler, O. Kupferman & M.Y. Vardi (2001): *Coverage Metrics for Temporal Logic Model Checking*. In: *Tools and algorithms for the construction and analysis of systems*, *Lecture Notes in Computer Science* 2031, Springer-Verlag, pp. 528 – 542, doi:10.1007/3-540-45319-9_36.

[9] H. Chockler, O. Kupferman & M.Y. Vardi (2003): *Coverage Metrics for Formal Verification*. In: *Correct Hardware Design and Verification Methods (CHARME)*, *Lecture Notes in Computer Science* 2860, Springer-Verlag, pp. 111–125, doi:10.1007/978-3-540-39724-3_11.

[10] Hana Chockler, Norman E. Fenton, Jeroen Keppens & David A. Lagnado (2015): *Causal analysis for attributing responsibility in legal cases*. In: *Proceedings of the 15th International Conference on Artificial Intelligence and Law, ICAIL*, ACM, pp. 33–42, doi:10.1145/2746090.2746102.

[11] Hana Chockler, Joseph Y. Halpern & Orna Kupferman (2008): *What causes a system to satisfy a specification? ACM Trans. Comput. Log.* 9(3), doi:10.1145/1352582.1352588.

[12] E. M. Clarke, O. Grumberg & D. A. Peled (1999): *Model Checking*. MIT Press, Cambridge, Mass.

[13] E.M. Clarke, O. Grumberg, K.L. McMillan & X. Zhao (1995): *Efficient generation of counterexamples and witnesses in symbolic model checking*. In: *Proc. 32nd Design Automation Conference*, IEEE Computer Society, pp. 427–432, doi:10.1145/217474.217565.

[14] N. Hall (2004): *Two concepts of causation*. In J. Collins, N. Hall & L. A. Paul, editors: *Causation and Counterfactuals*, MIT Press, Cambridge, Mass.

[15] J. Y. Halpern & J. Pearl (2005): *Causes and explanations: A structural-model approach. Part I: Causes*. *British Journal for Philosophy of Science* 56(4), pp. 843–887, doi:10.1093/bjps/axi148.

[16] Joseph Y. Halpern (2015): *A Modification of the Halpern-Pearl Definition of Causality*. In: *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI*, AAAI Press, pp. 3022–3033.

[17] J.Y. Halpern & J. Pearl (2001): *Causes and Explanations: A Structural-Model Approach — Part 1: Causes*. In: *Uncertainty in Artificial Intelligence: Proceedings of the 17th Conference (UAI-2001)*, Morgan Kaufmann Publishers, San Francisco, CA, pp. 194–202.

[18] Y. Hoskote, T. Kam, P.-H Ho & X. Zhao (1999): *Coverage estimation for symbolic model checking*. In: *Proc. 36th Design automation conference*, pp. 300–305.

[19] D. Hume (1939): *A treatise of human nature*. John Noon, London.

[20] O. Kupferman (2006): *Sanity Checks in Formal Verification*. In: *Proc. 17th International Conference on Concurrency Theory*, *Lecture Notes in Computer Science* 4137, Springer-Verlag, pp. 37–51, doi:10.1007/11817949_3.

[21] O. Kupferman & M.Y. Vardi (2003): *Vacuity detection in temporal model checking*. *Journal on Software Tools For Technology Transfer* 4(2), pp. 224–233, doi:10.1007/s100090100062.

[22] D.A. Lagnado, T. Gerstenberg & R. Zultan (2013): *Causal responsibility and counterfactuals*. *Cognitive Science* 37, pp. 1036–1073, doi:10.1111/cogs.12054.

[23] R.A. De Millo, R.J. Lipton & F.G. Sayward (1978): *Hints on test data selection: Help for the practicing programmer*. *IEEE Computer* 11(4), pp. 34–41, doi:10.1109/C-M.1978.218136.

[24] R.A. De Millo & A.J. Offutt (1991): *Constraint-based automatic test data generation*. *IEEE Transactions on Software Engineering* 17(9), pp. 900–910, doi:10.1109/32.92910.

[25] *RuleBase PE Homepage*. http://www.haifa.il.ibm.com/projects/verification/RB_Homepage.

[26] T. Schubert: *High level formal verification of next-generation microprocessors*. In: *DAC'03*.

[27] C.-J. H. Seger & R. E. Bryant (1995): *Formal Verification by Symbolic Evaluation of Partially-Ordered Trajectories*. *Formal Methods in System Design* 6(2), doi:10.1007/BF01383966.

[28] C.-J. H. Seger, R. B. Jones, J. W. O'Leary, T. F. Melham, M. Aagaard, C. Barrett & D. Syme (2005): *An Industrially Effective Environment for Formal Hardware Verification*. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems* 24(9), doi:10.1109/TCAD.2005.850814.

[29] S. Tasiran & K. Keutzer (2001): *Coverage Metrics for Functional Validation of Hardware Designs*. *IEEE Design and Test of Computers* 18(4), pp. 36–45, doi:10.1109/54.936247.

[30] J. Yang & A. Goel (2002): *GSTE through a case study*. In: *ICCAD*, doi:10.1145/774572.774651.